

CONFIDENTIAL DATA

Computer Experiments in Plasma Physics

by

O. Buneman and D. A. Dunn

GPO PRICE \$ _____

CFSTI PRICE(S) \$ _____

Hard copy (HC) 3.00

Microfiche (MF) 1.75

ff 653 July 65

January 1965

Report No. SU-0254-2

Prepared under

National Aeronautics and Space Administration

Contract No. NSG 299-63

N66 29435

(ACCESSION NUMBER)

(THRU)

(PAGES)

(CODE)

(NADA OR OR TRM OR AD NUMBER)

(CATEGORY)



INSTITUTE FOR PLASMA RESEARCH
STANFORD UNIVERSITY, STANFORD, CALIFORNIA

COMPUTER EXPERIMENTS IN PLASMA PHYSICS

by

O. Buneman and D. A. Dunn

January 1965

*Reproduction in whole or in part
is permitted for any purpose of
the United States Government.*

Report No. SU-0254-2

Prepared under
National Aeronautics and Space Administration
Contract No. NSG 299-63

[REDACTED]

Institute for Plasma Research
Stanford University Stanford, California

ABSTRACT

29435

The equations of motion of many interacting particles can be solved simultaneously using modern computers. Given the initial conditions, the subsequent behavior of all of the particles can be calculated and examined in graphs and motion pictures. The results of such computer experiments have led to new insights in many areas of plasma physics.

Plasmas are partially or fully ionized gases. A gas may be ionized by simply heating it or a plasma may be produced synthetically by combining electrons and ions from sources of these particles. A plasma is nearly electrically neutral; there are nearly equal numbers of negatively charged electrons and positively charged ions, but there may be large relative drift motions of one species, such as the electrons, through the ions within the nearly neutral plasma.

Plasmas were studied intensively in the late 1920's and early 1930's and then work in this field tapered off until after the hydrogen bomb was exploded. Then it became clear that controlled thermonuclear fusion reactions of the light elements were at least theoretically possible, if very hot plasmas of very high densities could be contained for times of the order of microseconds to milliseconds. A very substantial research effort in this field has since been carried out in the U.S., the Soviet Union, England, France, and West Germany. In the U.S. the allocation of funds to this program has been at the rate of about 25 million dollars per year and the British effort has been at about half this level. The payoff that is hoped for is a new source of power in almost unlimited amounts with near zero fuel costs.

One clear result of about 7 years of hard work is that a controlled thermonuclear reactor is not easy to build. The most pressing need in this area has been a deeper understanding of the behavior of plasmas. Time after time machines have been built that failed to function as expected, and then in understanding why the machine failed, a new type of plasma behavior has been uncovered. There are many other applications for plasmas such as ion propulsion, magnetohydrodynamic power generation, thermionic energy converters, and high-power, high-speed switches. In addition the fields of astrophysics and of geophysics are largely concerned with the behavior of the plasmas that fill most of the space beyond the atmosphere. Outside of a few planets, the universe is believed to be in the plasma state of matter, both in interstellar space and within the stars themselves. In all of these areas the basic static behavior of plasmas is just as could have been predicted in the 30's from the ideas and experiments of Irving Langmuir and his co-workers. The

difficulties arise in connection with the dynamic, time-varying behavior of a plasma that is not in thermodynamic equilibrium.

The physical laws that govern most of the dynamic behavior of plasmas were known nearly 100 years ago--they are Newton's laws of particle motion and Maxwell's field equations. Yet plasmas are poorly understood media. They have been called capricious and unpredictable. There is general agreement that such present-century discoveries as relativity and quantum mechanics have little or no bearing on most of the basic mysteries of plasma phenomena, nor can the discovery of new nuclear interactions and strange particles do anything to solve the problems of the plasma physicist. The simple reason why we have failed to understand plasmas is that they contain too many simultaneously interacting particles. The law of interaction between any pair of particles is, however, known accurately.

The behavior of a large population of interacting particles, persons, molecules, or other units can be approached from two rather different points of view. In the first method, the laws governing the behavior of the individual particles when interacting with one or more other particles are carefully described for each particle. Then a set of equations, with typically one or two equations per particle, is set up to describe the overall behavior of the total population and solved simultaneously.

In the second method, the system of interacting particles is viewed as a single aggregate which has properties as an aggregate that may vary with position and time, but that are not dependent on the individual particle behavior, except in a statistical sense. Normally this second approach leads to a much simpler analysis, requiring only a few equations to describe the entire population, and it has formed the basis for much of the progress in physics, as well as in many other scientific disciplines.

In physics, the individual particle model was first applied to celestial mechanics, where many interesting problems involving the motion of small numbers of particles existed. The approach was generalized and developed to an advanced level by Joseph Lagrange (1736-1813) and is generally referred to as the "Lagrangian" model. The second model has been used to deal with the motion of fluids and gases, both in the conventional mechanics of uncharged media and in plasma physics where charge

and current distributions are important. Leonhard Euler (1707-1783) was responsible for bringing the fluid model to a high state of development and one speaks of the "Eulerian" model in this connection. In the Eulerian approach, one sits and counts the particles present at each instant in every small region of space and that have velocity components within certain small limits. In the Lagrangian approach one attributes individuality to the particles and then follows each particle through space and time. One takes a census of the population only when needed for the evaluation of macroscopic fields, or when macroscopic averages are definitely asked for.

The Eulerian model has been used primarily for the study of near-equilibrium conditions in a medium, so that a state of statistical equilibrium can be used as a starting point and small perturbations from equilibrium in the velocity distribution or particle density can be studied. The solution of plasma problems with statistical thermodynamical methods has proved difficult in general because, in highly ionized media, all particles are interacting at all times, not just during close collisions. In the statistics of ordinary non-ionized gases the particles can be treated as if most of the time they do not interact. However, even when the difficulties of long-range interaction are overcome, a wide variety of important problems remain that are not easily handled by the statistical Eulerian method, because the most interesting and important conditions are those which are not near equilibrium. A typical example is a transition from an orderly state to a turbulent or disorderly state.

Since the statistical approach to plasma dynamics is so limited, we may ask if the Lagrangian approach can be applied directly, using one or two equations per plasma particle and a very large computer to solve these equations simultaneously. This scheme seems ambitious. In principle, one ought to program the dynamical equations for, say, 10^{15} plasma particles involved in a typical plasma phenomenon and solve them step-by-step in time just the way nature does it. In practice, one can build a sufficiently realistic coarse-grained model of the plasma in the computer, containing typically between 10^3 and 10^4 particles. Fortunately, computers are just getting large enough and fast enough to handle such a

problem, and fortunately, significant and reliable results can be obtained from such a coarse-grained model. We can make it behave outwardly just like a plasma which one sets up in the laboratory.

This half-way house between elegant theory and experimental hardware, our programmed version of the physical laws and boundary conditions, we call a "computer experiment." It differs from a typical computation of a theoretical result in that we do not evaluate mathematical expressions derived from the laws of nature, but we make the computer simulate the physical system. It is an analog of the physical plasma, although in most cases a digital computer is used in the simulation.

There are several important advantages of the computer experiment over the actual experiment:

1. it can be made to go in slow motion,
2. any quantity of interest can be observed without interfering with the experiment: probes, thermometers, voltmeters, etc. can be stuck into the simulated plasma at any position,
3. individual and collective particle behavior can be watched and displayed in a variety of ways--graphs, monitored records, oscillograms, and motion pictures,
4. boundary conditions can be controlled, space situations can be simulated,
5. special processes (ionization, optical excitation, hard collisions) can be switched on and off at will and their effect on the system behavior explored separately.

The role of statistics is somewhat peculiar and novel in these computer experiments. In the first place, it should be understood that the name "Monte Carlo" method would be somewhat misleading if applied to these calculations. We often solve strictly deterministic problems, with no element of chance introduced anywhere, and a close watch is kept on the randomizing effect of rounding-off errors in the computer. Indeed, Dawson has, in one of his computer plasma experiments at Princeton University, "turned the clock back" and reobtained the initial state of his plasma quite accurately in a strictly deterministic run.

At times, however, computer economy requires us to throw away information, use crude but fast integration procedures, coarse-grain and thus roughen-up the smooth and rigid course of deterministic classical physics.

Elements of chance are thus introduced which are not really wanted--although there are situations where such elements of chance impart a dash of realism!

We do not deliberately use statistical procedures in order to abstract from them some smooth and steady data, as one would do in the "Monte Carlo" method. A deliberate programming of statistical laws occurs only at boundaries--hot surfaces which emit randomly distributed particles and whose internal dynamics could not be programmed into the computer.

Order-disorder transitions are, perhaps, the most interesting results of our calculations. We have started systems cold, regular and well ordered. In a short time they have evolved into such complexity that the human mind can only comprehend them in a statistical manner (see Fig. 11). The computer may, in the course of this evolution, not have discarded any information--it has not increased the entropy--but the observer, absorbing only an incomplete averaged picture, has in his mind increased the entropy by discarding information!

Again, we have observed at times that there can be two kinds of fluctuations superimposed on an ideal, steady, clear average. One of these could be traced to the discrete and random nature of the particle injection procedures which were programmed at boundaries. These fluctuations come under the heading "shot noise." The other type of fluctuation is more violent. It would occur even if the boundary conditions were kept perfectly ordered and continuous. This is turbulence: the system becomes unstable, as does laminar flow under certain conditions, and directed initial energy is divided up into smaller and smaller packets, becoming random eventually (see Figs. 13, 18, and 20).

It is with a view to the latter possibility, into which we have run more often than expected, that our computations are kept evolutionary or sequential. We follow the system in time, we do not assume the existence of a steady state. The method is not, like several established methods of dealing with charged particles in electromagnetic fields, iterative. We are not trying to home on an elusive and possibly non-existent steady state. If such a state exists, and if it is stable, our system will converge upon it. If the steady state is only a calculable time-average buried under large fluctuations, we shall recognize it as such.

The sequential procedure is ideally suited to digital computers which, after all, work in a time sequence themselves--even when they have to solve "simultaneous" equations. In tracing particles through their mutual fields, we calculate the fields at every time step. Since time is kept as a real variable, and since we do not from the start assume our solutions to oscillate or to grow or decay exponentially in time, our transient analysis is not restricted to linearized ideal systems. Non-linearities present no problems. "Wave-wave" interactions are built in. Indeed it is the nonlinearity of a plasma which often hides the most interesting and hitherto least understood phenomena that the computer experiment reveals.

NON-INTERACTING PARTICLE EXPERIMENTS

A particularly simple Lagrangian problem is that of a single particle moving in a given electric and magnetic field. A good estimate of the behavior of some types of many-particle systems can be obtained by assuming that the particles are non-interacting. Many single particle solutions for different times or positions can then be superimposed to get the approximate many-particle solution.

A plasma problem of considerable interest in thermonuclear research is the problem of many charged particles moving in a magnetic mirror field of the type shown in Fig. 1. As a first step in understanding the behavior of such a system it is appropriate to study the behavior of a single particle in a magnetic mirror. This is a very simple problem at first glance. However, even when the magnetic field shape is axially symmetric, as in Fig. 1, and can be described by a simple equation such as a sinusoidal function, it has not been possible to get a satisfactory analytic solution to the motion of a single particle in this field. Störmer devoted a lifetime to tracing orbits in the simple dipole magnetic field of the Earth. His results have been useful in the understanding of cosmic ray fluxes and Van Allen-belts. In order to get an answer to this type of problem analytically, it is necessary to use an adiabatic




FIG. 1. A magnetic mirror is an important containment system for hot plasmas, but the exact solution for the motion of only a single particle in such a system is beyond present mathematical techniques. An approximate theory predicts that particles with initial velocities lying within the "adiabatic loss cone" will escape.

approximation to the trajectory, and this implies that there are only very small changes in any single loop about the flux line. It turns out that many quite ordinary trajectories are not very well described by this adiabatic theory, and a computer experiment turns out to be a useful approach to this problem. One of the authors¹ and Mr. R. E. Holaday studied this problem at Stanford University for the case of particles injected along or nearly along a flux line using an analog computer. This case is not well suited to the adiabatic theory, because the first turn of the particle around a flux line is always a large perturbation of its previous path. Figs. 2a, b, and c show three typical trajectories for this case for different angles of injection. For the case shown, the

Trajectories in the r - z plane obtained from a computer experiment for a single particle in a magnetic mirror. In all three cases the magnetic flux lines are the same, with one line about which the particle spirals that starts parallel to the ζ -axis at $\zeta = 0, \rho = 1.0$ where the particle is injected and ends up again parallel to the ζ -axis at $\zeta = \pi$. For this case the magnetic flux density is 10 times as great at $\zeta = \pi$ as it is at $\zeta = 0$.




FIG. 2a. Here the particle was injected along the flux line and the magnitude of the initial velocity was adjusted to just produce mirroring at $\zeta = \pi$. Adiabatic theory predicts that all such particles should escape from the mirror.




FIG. 2b. Here the particle had the same initial speed as in Fig. 2a, but its velocity was directed outward at an angle to the flux line. Mirroring occurred nearer to the injection plane.




FIG. 2c. Here the particle had the same initial speed as in Figs. 2a and 2b, but its velocity was directed inward. Mirroring did not occur. Evidently the exact loss cone is tilted inward and depends on parameters other than those of adiabatic theory.

particle with the inward velocity passes through the mirror, but the two other injection angles result in mirroring. The computer experiment predicts that the conditions for mirroring include the magnitude as well as the angle of the initial velocity with respect to a flux line and the distance off-axis that the particle starts. In the adiabatic approximation the condition for mirroring is completely determined by the angle that the initial velocity makes with the flux lines. Further computer experiments on this problem have been carried out by Roth at Cornell University.

Another area in which single particle computer experiments have proved useful has been that of microwave electron tubes. In many such tubes the electric field produced by the microwave circuit is much larger than the field acting on a particle due to other nearby particles. In this situation the motion of individual particles can be calculated separately and their orbits superimposed to obtain an estimate of the behavior of many particles. In the first paper on velocity modulation, the principle by which the klystron functions, O. Heil and A. A. Heil performed a large-signal analysis using the Lagrangian model with non-interacting particles, as did D. L. Webster in a theory of the klystron published somewhat later. This type of non-interacting particle theory was applied to the traveling-wave tube by Nordsieck in 1949, and has been useful in many other microwave tube problems.

Figure 3 is a distance-time diagram for a klystron amplifier which shows the trajectories of many particles which, under the assumption that the particles do not interact, are all straight lines in this representation. The slopes of these lines are their velocities, and in the klystron the velocity of the particles is varied sinusoidally as a function of entry time by the voltage of the input resonator, also shown as a function of time. In this type of tube the velocity of each particle, i.e., the slope in Fig. 3, remains constant after entry. It is evident from this diagram how a sinusoidal variation in velocity at one position is converted to a variation in particle density as a function of time at a position such as that marked $X = 1$ or $X = 1.84$. In a klystron this particle density variation with time is used to drive an output resonator at a position such as $X = 1$ or 1.84 that provides microwave energy to a load.




FIG. 3. A distance vs. time diagram for a klystron amplifier (an Applegate diagram). Each line represents the path or trajectory of one particle. The particle velocity is the slope of a line in this diagram which shows the paths traveled by many different particles leaving the injection plane at $X = 0$ with velocities that vary sinusoidally with time. For this simple model the particles do not interact; they travel at constant velocity and so the line for each particle is straight and has a value determined by the initial velocity. At a position such as $X = 1$ there is strong bunching and an output cavity placed at this position would be driven by the charge variation in time (current) that exists at this position. (From K. R. Spangenberg, Vacuum Tubes, McGraw-Hill, 1948).

In microwave tube work, the distance-time diagram is often referred to as an Applegate diagram. L. M. Applegate was the patent attorney for R. H. Varian and W. W. Hansen, the inventors of the klystron, and he used this diagram to explain the operation of a klystron in the Varian-Hansen patent. The distance-time diagram is also widely used in connection with the motion of particles in relativity theory, and in this connection it is referred to as a Minkowski diagram. We will be using it extensively in the remainder of this article to describe the motion of many interacting particles.

INTERACTING PARTICLE EXPERIMENTS WITH AXIAL OR PLANAR SYMMETRY

The two and three-body problems of Keplerian celestial mechanics are very close to the modern problems of plasma physics involving many interacting particles. In both types of problem the particles interact with an inverse square-law force, although of course there is no direct analog of the repulsive force between like charges in gravitational problems. It is interesting that it was precisely the difficulty with the computations that slowed progress in this field at the time of Kepler, although conceptually the problem was clear with three or more interacting bodies. There is, as yet, no rigorous analytical solution to the general three-body problem, and it is still nearly impossible to handle large numbers of particles with charges of both positive and negative sign moving in three dimensions even with modern computers. However, considerable progress has been made by taking advantage of the symmetry inherent in many physical problems in plasma physics. For example, consider the cylindrical system of Fig. 4 in which two parallel plane surfaces oppose each other. We expect to be able to learn a great deal about such a system by studying modes of operation in which there are no variations in electron current or any of the other variables as we move around the system azimuthally, i.e., modes that are symmetrical about the axis. Still a further simplification is obtained by allowing no variations in the radial direction. This latter simplification is especially appropriate if the dimensions

of the system are such that the separation between the end surfaces is small in comparison with the diameter.




FIG. 4. A plasma diode with symmetry about an axis through the electrodes. The drawing can represent a dc gas discharge tube if electrode 1 is an emitter and electrode 2 is operated at a potential higher than electrode 1 so a current is drawn through the ionized gas or plasma consisting of electrons and ions plus neutral molecules. A thermionic energy converter can have a similar construction.

On the computer we might simulate the system by either a set of disc particles or the infinite sheet particles of Fig. 5. In both cases we imply particle motion in only one dimension, along the longitudinal axis, such as might occur if there were a strong axial magnetic field present. Alternatively, a physical system without a magnetic field might behave in the same manner, if the axial velocity of the particles were large in comparison with the transverse velocity and if the diode spacing were




FIG. 5. Lagrangian models of plasmas can be designed around charged particles that are point charges with an electric field that varies as a square law and that can move in three dimensions, rod charges with a $1/r$ field variation that can move in two dimensions, or plane or disc charges that can move in only one dimension with field variations as shown. There is a tradeoff between the number of particles that can be included in a problem and the number of dimensions in which motion is allowed.

small enough. For many geometries the disc particles would not behave significantly differently from the sheet particles. For a diode problem with the diameter comparable to or less than the axial spacing, however, some difference is to be expected. The difference in the two models is indicated in Fig. 5 where the Coulomb law for each of the two cases is shown. Naturally it costs more to do a problem with a force law that depends on distance than for the case where the force is independent of distance.

A further refinement, still for the case of one-dimensional motion, would allow a series of ring-shaped particles to slide past each other and travel through the diode at different velocities at different radii, but with each ring remaining fixed in size. Obviously subdividing the disc particles into rings requires additional computer storage. For many problems it is also interesting to allow motion transverse to the central axis. In cylindrical geometry this two-dimensional motion problem would be simulated by ring-shaped particles that could expand and contract. This, again, raises the demand on storage space, quite apart from the increase in integration time and complications in field evaluation. We can also have two-dimensional motion between parallel planes and here the particle shape that corresponds to the expandable ring in axial symmetry is an infinitely long rod. Considerable work has been done with this two-dimensional rod-electron model, and many of the results can be carried over qualitatively to cylindrical problems.

In order to illustrate these ideas more specifically, let us consider Fig. 6 which is a schematic drawing of a diode with five charge sheets in one-dimensional motion at arbitrary instantaneous positions within the diode. The distance across the diode is normalized to unity and the charge sheets are numbered from right to left, $i = 1, 2, \dots, 5$. At a given time step let us imagine that the charge sheets have the positions x_1, x_2, \dots, x_5 indicated by the solid lines. The normalized electric field within the diode is given by the equation given in the figure and is also shown below the diode. The field is seen to step by one unit each time a sheet is passed. The electrostatic potential corresponding to this field




FIG. 6

FIG. 6. A diode with both walls at the same potential containing five plane charge sheets (solid lines) all of the same (positive) charge. A calculation of the electric field for this case is extremely simple and is given by the equation $E_j = \frac{1}{2} + j + \sum_{i=1}^k x_i$ at the position of the j th charge sheet. Between the charge sheets to the right of the j th sheet the field is $E_{rj} = 1 - j + \sum_{i=1}^k x_i$ and to the left of the j th sheet $E_{lj} = -j + \sum_{i=1}^k x_i$. For this case $\sum_{i=1}^k x_i = 2.5$. The potential V is the integral of the electric field. In this figure x is measured from left to right and charge sheets are numbered from right to left. The dashed lines suggest possible positions for the charge sheets after they have been allowed to move under the influence of the field given here for one time step. The positions of the charges after one time step also depend on their velocities at the beginning of the time step as well as their positions and the field.

is also shown in Fig. 6 for the case of positively charged sheets and with both diode walls at the same potential.

It is evident what a simple calculation is required to determine the field and potential from the positions of the sheets. Evaluation of the field consists primarily of counting the number of sheets to one side of any position. For the potential, one simply computes the sum of all the positions and then, moving across the diode, subtracts the number of the sheets passed at any point.

After making this computation, the charges are allowed to move under the influence of the local field for the duration of one time-step, starting from their known positions with their known velocities. New positions, suggested in Fig. 6 by the dashed lines, are then obtained and the field is recomputed and the process repeated.

Each of the above steps can be written as a set of finite-difference equations with the differencing interval being an important computational parameter that must be made so small that the results of the computer experiment are unaffected by making it smaller. The basic process that we are using in these computer models is one of coarse-graining. Not only must the time-step be made short enough, but also the number of computer particles must be large enough to properly represent the phenomenon being studied. There is usually a significant wavelength in the problem and typically there must be of the order of 20 particles per wavelength to give a result that is invariant to an increase in the number of particles. The coarse-graining must not be too coarse, either in time or particles per wavelength. Sometimes the wavelength is totally unknown when the computer experiment is begun, and then, as in the case of the time-step, the number of particles must be increased in successive experiments until the results do not change when the number is further increased.

The basic computational process in more complex cases than the one-dimensional diode is still the same in principle although more involved in detail. It consists of two steps that are recycled time-step after time-step:

1. Compute the field from the charge positions.
2. Compute new positions for all of the charges after they have been acted on for one time-step by the field computed in Step 1.

Figure 7 is a diagram of this basic process. If there are charges of more than one sign, self- or applied-magnetic field effects, two-dimensional electric or magnetic field effects, or two-dimensional motion, the process is more complex in either or both steps than for the problem suggested in Fig. 6.

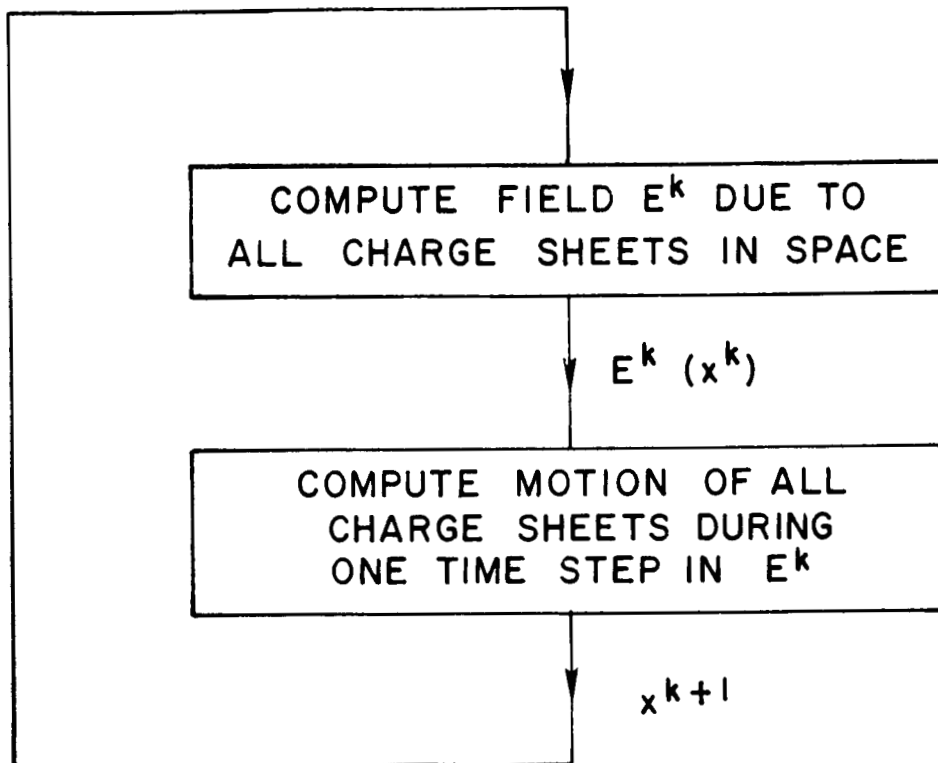


FIG. 7. The basic computational process is diagramed. E^k is the field at the k th time step which is computed from the charge positions x^k at that time step in the manner suggested in Fig. 6. The motion of the charge sheets in this field is computed and a new set of positions x^{k+1} is found which determines a new field pattern $E^{k+1}(x^{k+1})$, etc.

All of these cases are greatly simplified as compared with a three-dimensional motion, three-dimensional field problem. The interesting and important feature of these simplified models is the fact that computations can be carried out on them for thousands of charge sheets in one-dimensional motion or hundreds of charged rods in two-dimensional motion in about the same time it would take to do similar computations on only a few charges in three-dimensional motion. A tradeoff between the number of dimensions of motion and the number of charges exists that can be highly advantageous to the computer experimenter, because there are many problems with great practical significance that can be simulated realistically by a few thousand charges. When more than one dimension is essential to an understanding of the phenomenon, a two-dimensional computer experiment with up to 2000 charged rods or rings is within the capability of present computers such as the IBM 7094 and interesting results have been obtained in such problems in about a half-hour of computing time. Problems involving large numbers of particles in three-dimensional motion must await the next generation of computers.

SINGLE-SPECIES DIODE EXPERIMENTS

Historically, it appears that Hartree and Nicolson were the first to take advantage of the symmetry tradeoff to make calculations on a many-body problem in which the particles were restricted to one-dimensional motion. They analyzed a planar electron diode using a one-dimensional sheet model in 1943 using a desk calculator. They used about 30 interacting particles and were able to study the buildup of circulating current in a magnetron. Their charge sheets experienced not only a force due to the electric fields of the other charge sheets plus any applied electric field, but also a magnetic force due to an applied magnetic field directed perpendicular to the direction of charge motion. They were able to show that the unexpected dc current flowing in a magnetron at magnetic fields well above a theoretical cutoff value was not a result of some transient

effect present in a one-dimensional model. Theories and computer experiments taking into account two-dimensional effects have since explained the observed phenomena. They were also able to show that a single-stream dc state predicted by Brillouin was the natural dc state and would be set up even when the magnetron was suddenly turned on with no charge in the space initially. The state breaks up, however, due to two-dimensional instabilities.

Results of a typical calculation made by Hartree and Nicolson are given in Fig. 8 which is a plot of the particle trajectories in the distance-time plane along with the voltage and current as a function of time. In Fig. 8 it is easy to see the qualitative result of the computer experiment, which is that the electron layers peel off the cathode successively, stay in order and move out to a certain position about which they oscillate after the voltage reaches its full value.

Some of the most interesting and important recent applications of computer experiments have been in electron and ion diode problems with either no magnetic field or a very strong magnetic field in the direction of current flow, as in Fig. 4, so that either the sheet or disc model of Fig. 5 can be used. Hartree² made a study of this type of problem in 1950 in order to determine whether it is possible during the transient turn-on period of an electron diode for electrons to reach the anode at a higher energy than that corresponding to the applied voltage. In his problem the particles were all injected into the diode at the same initial velocity and the two walls of the diode were maintained at the same potential. The diode was initially empty. Hartree carried his computations only a short time beyond the point where the diode fills completely and observed no unexpected behavior.

Unfortunately Hartree died before he was able to see the result obtained by R. J. Lomax³ who had, at his suggestion, carried the experiment further and obtained a very intriguing and fundamental effect that appears at high currents after the diode fills. Birdsall and Bridges⁴ observed it in Berkeley at about the same time as Lomax found it in Cambridge, England. The effect is illustrated in Fig. 9 which again shows a set of electron




FIG. 8. Distance vs. time trajectories of charge sheets in a one-dimensional diode in the presence of a transverse magnetic field. Each charge sheet has zero initial velocity. The accompanying current and voltage curves are shown as a function of time. The voltage is specified as the dashed line and the plate separation is sufficiently large that none of the charge sheets reaches the opposite side, so there is zero steady-state current, simulating a magnetron beyond cutoff. (After Hartree and Nicolson, 1943.)

sheet trajectories in the distance-time plane. In this case all the electrons are injected at the same velocity at a uniform rate and the electron diode is infinite in length, i.e., the anode is very far away. Hence, all the injected current must return to the injection plane. An attempt to eject charges of a single sign from a space ship would have this result. Thus an ion propulsion system must include means for neutralizing the ion beam with charges of the opposite sign. The static theory for this case predicts that all the electrons turn around at the same place, at $X = 0.707$ in Fig. 9, and that there are no oscillations. The result of the computer experiment is that the electrons turn around at different positions and that the position for turnaround varies with time, as does the current returned to the injection plane. It turns out to be impossible to set up the static state starting with an initially empty diode and, if it is artificially created as an initial condition, it immediately is transformed to the state of Fig. 9 which is the natural mode of the system. This natural mode is seen to be an oscillatory one. The result is the same for finite length diodes for high enough values of current.

For many years the static states of an electron diode had been assumed to be the only solutions to the equations describing such problems. However, it was clear from the static theory that there was an upper limiting current that a given diode could transmit. It was assumed by the early workers in this field that, if a current greater than the limiting value were injected, the system would switch to another static state with some current being returned to the injection plane. Diagrams tracing out hypothetical transitions from one static state to another are shown in many textbooks on electron tubes such as the standard works by A. H. W. Beck and K. R. Spangenberg, but this behavior has never been measured. In fact the system does not do as postulated. Instead it adopts a time-varying solution to the problem with quite large amplitude time variations as in Fig. 9. This time-varying state is, in many cases, an oscillation around a mean at some definite frequency, and the mean is often an analytically calculable static state.




FIG. 9. Distance vs. time trajectories of electron charge sheets in an infinitely large one-dimensional diode injected with an initial uniform velocity. All of the charge sheets return to the injection plane at $X = 0$, but their trajectories oscillate in the depth of penetration into the space and the number of charges returning to the injection plane per unit time (the current) oscillates in time. The frequency of oscillation is the plasma frequency corresponding to the average charge density at the turn-around point. A static, non-time-varying solution to this problem can be obtained, but it is not a stable one. It, however, correctly predicts the time-average turn-around point $X = 0.707$ for this problem.

In the electron diode above limiting current, the oscillation conditions can be usefully related to the static states, but the oscillating state is not an alternation between static states, nor is it a "relaxation" oscillation. The current to either electrode varies nearly sinusoidally in this case, and the frequency is near the electron "plasma frequency" of the part of the region with the highest charge density. No ions are present in this case and the "plasma oscillations" are those of the negative charges of the electrons at their characteristic frequency.

The plasma frequency is the natural frequency of oscillation of a charged medium and is proportional to the charge density. A plasma oscillation is the result of the electrostatic repulsion with which a compression of like charges is opposed. The electron sheet model (Fig. 6), with some uniform stationary neutralizing ion background, readily exhibits and explains this behavior, which prompted J. M. Dawson⁵ of Princeton University to simulate such a set of neutralized sheets by a series of pendula. For a typical electron diode the natural frequency might be of the order of 10^8 cycles per second. For an ion diode the static and dynamic theory is identical to that for an electron diode, but the oscillation frequency is lower because of the heavier mass of an ion. The plasma frequency is inversely proportional to the square root of the particle mass and so ions in a cesium diode, without electrons, would have a natural oscillation frequency of the order of 10^6 cycles per second, if it were operated beyond limiting current. An oscillation at about this frequency has in fact been observed by J. M. Sellen in some experiments at Space Technology Laboratories with an un-neutralized cesium ion beam and is believed to be of this type.

OSCILLATIONS IN THERMIONIC ENERGY CONVERTERS

A thermionic energy converter is simply a diode as in Fig. 5 with electrons and ions emitted from one side of the diode which is maintained at a high temperature by an external heat source. The energy given to

the electrons by this source of heat energy external to the diode, such as a nuclear source, is converted to usable electricity by the process of the electrons flowing as a current through an external load resistor. In experiments on such devices it is found that in the regimes of most interest the electron current that flows is time-varying and is not a purely dc current as predicted by the static theory. One of the authors and P. Burger⁶ at Stanford University have recently conducted a series of computer experiments on a diode of this type in which the initial velocities of the electrons and ions are chosen from a source of random numbers to provide a simulation of electron and ion temperatures corresponding to the temperature of the emitting electrode. One result is a set of current vs. time curves, a typical one being shown in Fig. 10. The




FIG. 10. A thermionic energy converter is a diode as in Fig. 4 with electrons and ions emitted from the left wall of the diode with random energies acquired from an external source of energy. Taking into account the work functions of the surfaces the potential within the diode increases from left to right when the converter is delivering electrical energy to a load. The potential within the diode varies with distance as shown at the right in this figure at several different values of time. The current also varies as a function of time in the regular pattern shown for an ion-rich case, which is the case of interest for practical applications. These curves were obtained from a one-dimensional computer experiment using about 10^4 ion and 10^4 electron sheets within the diode. The current vs. time curve agrees with experimentally measured waveforms.

current in Fig. 10 is seen to vary in a characteristic manner with a rapid variation near the peaks of the slow variation. This waveform is in very close agreement with what has been measured by a number of experimenters for operating conditions similar to those simulated in the computer experiment.

By observing the voltage vs. distance across the diode as a function of time it has been possible to arrive at a qualitative explanation for the oscillations. The key to an understanding is the realization that the electrons and the ions oscillate with characteristic frequencies that are widely different. The ions are responsible for the slow variation in Fig. 10 and the electrons for the superimposed fast variation. If we now imagine the ions in a given pattern at some instant of time, the electrons can change their density distribution very rapidly, before the ions can respond. If we find that the electrons can have an alternative distribution of potential and density according to the static theory, temporarily regarding the ions as having infinite mass, we may find that the electrons will redistribute themselves and either select this alternative distribution or oscillate with a time-average density distribution that is different from that required by the original static state. When this happens, the ions will see a time-average force acting to disturb their equilibrium, and, being actually of finite mass, will tend to follow the electrons and hence to start an oscillation at the characteristic ion frequency. This process is also illustrated in Fig. 10 which shows a set of "stills" selected from an actual movie sequence. The calculations were carried out on a computer that provides a direct output on an oscillograph and the potential variation from plate to plate was projected on the display tube at each instant and photographed. One was thus able to watch the diode perform in slow motion and to draw general conclusions regarding its mode of operation. The state with the low, more or less constant potential region at time $T = 30.4$ in Fig. 10 is the one expected from static theory with finite mass ions. The higher potential state at time 31.2 is a possible one from static theory for which the ions have the same distribution of density as for the other state, but for which the electrons

are moving more rapidly and hence have a lower density. In the computer experiment it was found that the low potential state was established first and the electrons then moved toward the high potential state, with the ions following them after a number of electron oscillation periods. The system then refilled with ions and recycled.

It is seen that in this fairly complex situation the computer experiment leads to results that can be interpreted with the aid of static theory, so the two approaches go hand-in-hand in leading us to a complete understanding of the process involved. The computer was also used to test modes of operation in which only one static state was theoretically possible under the assumed conditions of temperature and electron and ion current. In such a case the computer rapidly homed on the static state, and, with the exception of small shot-noise fluctuations, stayed there. In fact, the computer homed on the state in a shorter time than it had previously taken to evaluate such a state from analytical formulae (involving error integrals). The computer was also used to test stability when two or more static states were predicted and the system was started in one such state. In such a case it was found to stay only temporarily in the selected static state before assuming its oscillatory condition which was the natural state. What is of real significance is that the computer runs preceded the development of the qualitative explanation. It was a case of "hindsight" allowing us to explain the computer results in terms of the possible static states of the system that were discovered to be involved.

It will be noted that the ion-to-electron mass ratio for the case illustrated in Fig. 10 is 16, a convenient value for computations. What is important is that the mass ratio should be much greater than unity in order to give results in qualitative agreement with physical experiments where the mass ratio may be several thousand. Again it is necessary to establish experimentally that the results do not change qualitatively as the mass ratio is further increased. It usually turns out that a mass ratio of 16 is sufficient to give good results. The reason it is

important to be able to use low mass ratios is that the computation time to go through a given number of cycles of the oscillation is normally proportional to the square root of the mass ratio.

TWO-STREAM INSTABILITY

If an electron beam is injected into a plasma or if an electric field is applied to a plasma, a stream of particles with a velocity very different from that of the rest of the plasma can be established. This particle stream will ultimately be scattered and randomized and the entire plasma will return to equilibrium, if the injected beam or applied field is shut off or if we go a great distance from the source of the field or the beam injection point in a "steady-state" situation. The most obvious processes leading to randomization are scattering from neutrals and scattering from other charged particles by means of close Coulomb collisions, i.e., a charged particle encounter that substantially alters the trajectory of at least one of the colliding particles. However, it was found experimentally in the machines that were built to study thermonuclear plasma containment and related problems that randomization occurs many times more rapidly than could be explained by close collisions. There has also been a series of experiments on dc gas discharges in which large amplitude high frequency oscillations at a plane very near the cathode of the discharge were observed, along with violent transverse scattering of the electrons beyond this plane. The electrons reaching this plane from the cathode are known to be in the form of a nearly single-velocity stream, while only a short distance beyond this plane in the region of the positive column, it has been well established that the electrons have a nearly perfectly random-velocity distribution. The mean free path for single electron collisions is hundreds or thousands of times the distance from the cathode to this oscillating plane and can't explain this rapid randomization.

This thermalization of an injected stream can be explained in terms of the collective interaction of large numbers of plasma particles

oscillating coherently. If one group of particles is moving with respect to another, this mutual collective interaction can extract energy from the moving stream and convert it to amplification of the collective oscillation. The basic process is known in the electron tube field as two-stream amplification, and it is a well-known technique for amplifying microwave signals by means of the interaction of two electron beams at different velocities. In the case of the plasma, one stream is normally standing still and the interaction can either be between an injected electron beam and the electrons of the plasma or between moving electrons and stationary ions. The case of electron-ion interaction as a small-signal amplification mechanism was also known in the electron tube field and was analyzed by J. R. Pierce in 1948. In the plasma thermalization process there is, of course, no applied microwave signal, and the amplification process starts from random noise in the plasma at frequencies near the frequency for maximum growth. Other frequencies are present and amplified, but to a lesser extent than the preferred frequency which soon suppresses all other signals.

What happens to limit the growth is illustrated by the trajectory diagram in the distance-time plane in an electron-ion interaction, as shown in Fig. 11. This figure is from a paper by one of the authors⁷ published in 1959. Similar computer work was done at Princeton University by Dawson who also published his work in 1959. In Fig. 11 the diagram represents one space-period out of a repeating series of identical diagrams that should be imagined connected on both sides of the one shown. Note that time runs vertically in this diagram. It is evident from Fig. 11 that after the large-signal limit is reached and the waves "break" they are very disorderly, if not entirely random. The net effect has been to convert the drift energy of the electron stream to a combination of high frequency field energy and random particle energy. This process is shown directly in Fig. 12 which is a plot of the space-average electron position as a function of time. The slope of this line is the electron drift velocity which is seen to be constant out to T of about 50 where it rapidly levels off to nearly zero velocity, at which point the drift energy has been converted to high frequency field energy and random energy.




FIG. 11a. Distance vs. time trajectories for electrons interacting with initially stationary ions with time plotted vertically. This diagram shows the electrons only. The ions are only slightly perturbed from vertical straight line trajectories during this time period and are shown in Fig. 11b. This diagram is to be understood as repeating to the left and right of the region shown; this is only one section out of an infinitely repeated pattern. An initial sinusoidal velocity modulation rapidly grows into an apparently random situation. Here the growth is in time rather than in distance.




FIG. 11b. Ion trajectories for the case of Fig. 11a.




FIG. 12. Average electron position as a function of time for the problem of Fig. 11. The time scale is the same as Fig. 11. Up to $T = 50$ there is a uniform average electron drift velocity (the slope of this curve). Beyond this time the electrons lose their drift velocity, converting their drift energy into random energy by the process illustrated in Fig. 11.

As a result of these calculations and graphic demonstrations of the effects of collective interaction, more realistic calculations and designs for thermonuclear experiments have been made that take this effect into account.

Similar large-signal calculations for microwave tubes have been very successful in predicting the efficiency and power output from various types of microwave amplifiers. The earliest work with interacting particles was done on traveling-wave amplifiers. Howard Poulter at Stanford University obtained a few results with the IBM 650 which was the fastest computer available in 1949, and when the IBM 704 became available an extensive series of calculations was made at Bell Telephone Laboratories by Tien, Walker, and Wolontis in 1955 and 1956. Further work has been done by J. E. Rowe at the University of Michigan and an extensive series of klystron calculations has been made by T. G. Mihran and S. E. Webber at the General Electric Company. In all of these calculations the results are similar to the two-stream results in Figs. 11 and 12 with the exception that the drift energy that is converted to high frequency field energy is removed as useful output power and the slowed-down thermalized beam is collected and lost as heat in the collector. Obviously in this application it is important to choose parameters that will give the maximum conversion of drift energy to field energy, and it is precisely this type of information that is uniquely supplied by the computer experiment.

ION PROPULSION

Fundamental in the use of ion beams for the propulsion of space ships is the requirement that the beam leaving the ship be neutralized. On the average, the same number of positive and negative charges must be shot out of the ship per unit time, i.e., there must be current neutrality which consists of equal positive and negative currents leaving the ship. Furthermore, these positive and negative charges must be ejected at nearly the same velocity to provide charge neutrality near the ship as well as current neutrality, so that none of the current will be returned to the ship. If

current of one sign is returned, the ship will charge up and attract back charges of the opposite sign, thus reducing thrust. The important question turns out to be, how different can the ejected velocities be and still provide charge neutrality?

Typical mission requirements within the solar system suggest the use of Cesium ions as a propellant at velocities corresponding to an energy of a few kilovolts. The simplest technique for neutralizing a beam of positive ions is the injection of an electron beam with the same current (charge per unit time) as the ion beam, thus forming a neutral plasma in the exhaust of the ship. If the two beams are ejected at the same velocity, there is no difficulty in predicting the result from static theory. In this case there is perfect charge neutrality throughout all space. However, it turns out that typical electron emitters, such as a hot tungsten wire, emit electrons with a mean random speed that is considerably higher than the velocity of Cesium ions at a few kilovolts. Thus, even if we don't accelerate the electrons at all before ejecting them from the ship, we find that we can't send them out with as low a mean velocity as the ions. What is more, they will inevitably have a large velocity spread in comparison with their drift velocity. This latter effect can be taken into account in the static theory, but the problem of a positive ion beam in the presence of an electron beam with a mean thermal speed greater than the ion velocity has a static theory solution that does not appear to correspond to what would be obtained in the physical situation. Indeed, H. Derfler at Stanford University has shown that there are conditions under which no static solution can be found at all and others under which all possible static solutions can immediately be proved unstable. Similar results are obtained with a single velocity electron beam that is traveling more than twice the speed of the ion beam. In the latter case, the static theory result requires the existence of a set of electron current loops that are not likely to be set up physically. The situation is reminiscent of the incorrect predictions of static theory for an electron diode operating above limiting current. We again wonder if there may not be a time-varying solution that is the correct description of the way the system would behave. When we perform computer experiments on this problem, we do indeed find that this is the case.

One of the authors⁸ and I. T. Ho performed some calculations at Stanford University on the problem of electrons and ions shot out of a space ship at different velocities, with both species having single velocities initially. For electron velocities less than twice the ion velocity the static theory results were reproduced by the computer model after an initial transient. However, for higher velocity electrons the results were different from static theory and, as in the electron diode beyond limiting current, the result was a time-varying one. Figure 9 shows what the results would be if only ions or electrons were ejected. They would return to the ship after forming an oscillating "sheath". Figure 13 shows some typical electron and ion trajectories for the case of equal electron and ion currents ejected from the ship. It turned out that for this case, when followed further in time, ions and electrons both returned to the ship in an oscillatory manner, so the ion beam was only partially neutralized.

However, after a little further experimentation it was found that complete neutralization could easily be obtained at the same or even higher electron velocities by the simple expedient of increasing the electron current. If this was done, any excess electron current beyond the necessary minimum was returned to the ship, and just enough electrons for perfect neutralization were supplied to the beam. Figure 14 shows electron and ion trajectories near the ship for such a case. It is seen that most of the electrons are slowed down nearly to zero velocity at a short distance from the ship and then some electrons are selected from this oscillating "sheath" to go on into space with the ions. The electron sheath is very close to the space ship and looks very similar to the pattern obtained without any ions in Fig. 9. As in Fig. 9 the oscillations here occur near the electron plasma frequency corresponding to the density near the sheath. In this case these oscillations continue out into space and the accompanying electric fields provide a mixing process for the electrons and ions that is a far more effective communication between particles than close collisions.

This result agrees with what has been measured in a number of experiments by J. M. Sellen and his co-workers at Space Technology Laboratories.




FIG. 13a. Distance vs. time trajectories for electrons ejected from a space-ship into an infinite space. In this case the electrons are accompanied by an equal injected ion current, and the electrons are shot into the space at a velocity 2.5 times the velocity of the ions. Only the electron trajectories starting out prior to $T = 52$ are drawn, although there was a continuous injection of current beyond this time. During this time period the ion trajectories were not significantly perturbed, and partial neutralization was obtained. However, at a later time some ions stopped and returned to the ship.




FIG. 13b. Ion trajectories for the case of Fig. 13a.




FIG. 14. Distance vs. time trajectories for electrons for the same case as Fig. 13 except that the injected electron current is twice the ion current. Some of the ion trajectories are shown dashed. For this case and for higher injected currents the ion beam is fully neutralized and does not re-turn to the space ship at all.

In these experiments it appeared that, once there was sufficient electron current available, the neutralization process was not critical and neutralization was obtained very easily. This result also agrees with qualitative pictures of the process that have been suggested by a number of workers in which the ion beam is characterized as a plasma bottle which is continually lengthening itself and within which the neutralization electrons are bouncing about.

Some very similar results are obtained when random initial electron velocities are put into the computer program. One of the authors⁹ has performed this one-dimensional computer experiment along with a number of other related experiments in cooperation with G. Kooyers and R. B. Wadhwa at Litton Industries. Figure 15 is a plot of the potential near the space ship for a one-dimensional case in which neutralization is obtained by means of electrons with a random Gaussian velocity distribution with a mean thermal speed greater than the ion velocity. This experiment included not only the region beyond the neutralizer grid but also the region between the ion gun and the neutralizer. The results of the computer experiment are drawn in Fig. 16 which is a plot of potential vs. distance as in Fig. 15, but plotted for a number of different values of time in a perspective view. The fluctuations of the potential in both time and space can be seen and compared with the height kT which represents the temperature of emitted electrons. The fluctuations are much greater than kT and again represent the result of a collective interaction

TWO-DIMENSIONAL EXPERIMENTS

In all of the work described so far, the motion of the charges has been in only one dimension. As indicated at the outset, two-dimensional motion can be included, but at a price in terms of the number of charges that can be followed. Figure 17 shows an ion beam neutralizer system consisting of electron emitters spaced apart by a small distance through which the ion beams can pass. Figure 18 shows the results of a calculation done by R. B. Wadhwa and one of the authors at Litton Industries




FIG. 15. A real ion propulsion system consists of an ion gun in which the ion beam is accelerated by an accel grid and a following hot decel grid at which electrons with random velocities are emitted. Electrons can and do travel in both directions from the hot decel grid neutralizer to form a neutral plasma that extends back nearly to the accel grid. This figure shows the potential vs. distance in such a system.




FIG. 16. Potential vs. distance for a number of values of time for a system of the type in Fig. 15. These results are from a computer experiment that included the electron temperature. The potential equivalent of the temperature kT is shown and it is seen that the fluctuations of potential in time and space are much greater than kT .

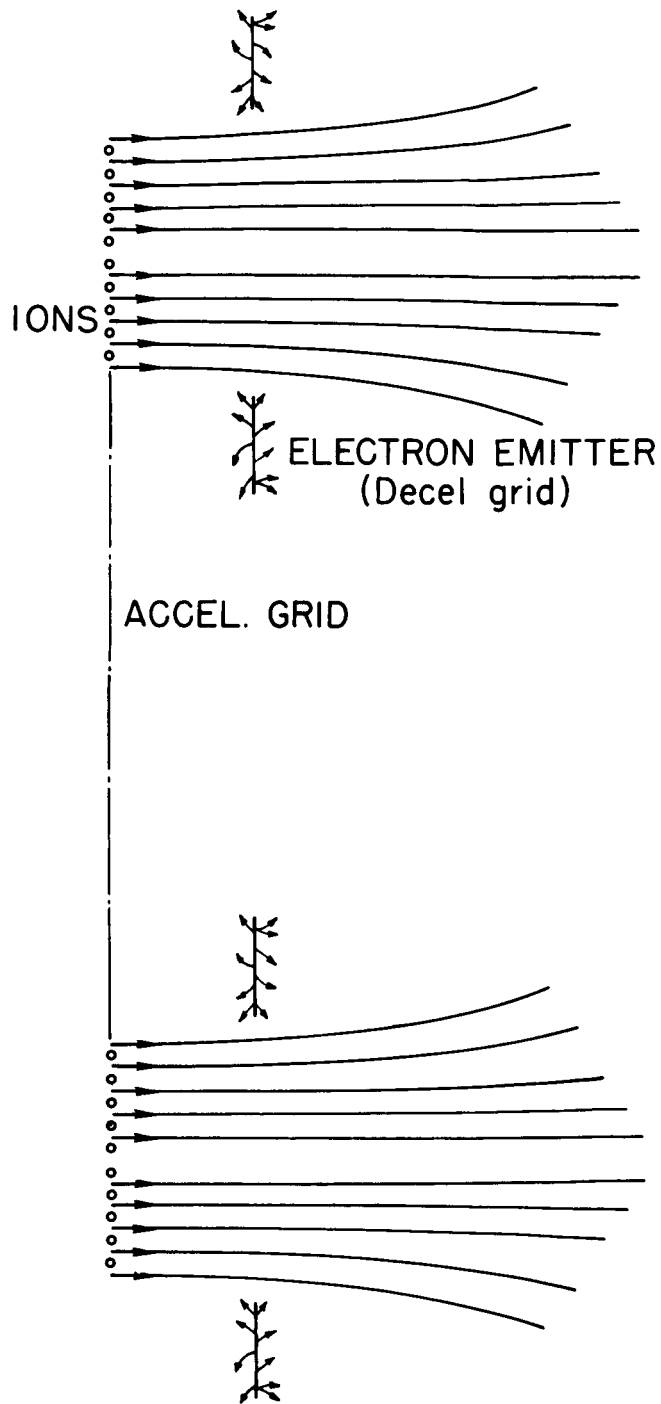


FIG. 17. An ion beam neutralizer system in which the electron emitter (decel grid) consists of a series of flat wires spaced apart so the ion beams can pass between the grid wires without interception. In this figure the system is infinite in the direction perpendicular to the paper and there is an infinite periodic array of ion beams and grid wires, of which we show only two periods.




FIG. 18a. Electron trajectories for the system of Fig. 17. Here only half of one period is shown. The pattern is symmetrical about the center of the ion beam, so this drawing tells the full story.




FIG. 18b. Ion trajectories for the case of Fig. 18 a.

including two-dimensional effects for both electrons and ions for such a system. The electrons in this case must not only slow themselves down to the right velocity and adjust their rate of flow correctly to provide neutralization, but also find their way from the emitters into the beams spaced transversely from the emitters. As seen in Fig. 18, this is exactly what the electrons do, and again the criterion for neutralization is simply that a sufficient excess electron current be provided. One conclusion from this beam neutralization work seems to be that nature loves a plasma to more or less the same extent that she abhors a vacuum.

In two dimensions, charged particles are treated as parallel rods. Since two coordinates and two velocities are required to describe each rod, the number of particles which a given memory can store is only half that for the one-dimensional sheet model. What is worse, however, is that these rods are now distributed over a two-dimensional area while the sheets were distributed along a single line only. Hence the linear resolution goes down considerably when passing from the one-dimensional to two-dimensional models.

In processing the rods through their dynamical evolution, i.e., in advancing each rod by one time step, the computing labor is doubled (in comparison with the sheet model), but there are no further complications. A magnetic field parallel to the axes of the rods is readily taken into account by coupling together the two components of motion: this case is important since the migration of charges across magnetic fields has often presented puzzling problems and mysteries in physical devices such as plasmas and electron tubes.

However, the real stumbling block in the two-dimensional problem is the calculation of the electric field that acts on each particle. A theoretician may be surprised that this step presents difficulties, for he knows the field equations to be linear while the dynamical equations in certain forms, are not. The digital computer, on the other hand, does not care or worry about linearity.

Several different methods have been devised for evaluating the fields. Perhaps the most direct method is to sum all the $N-1$ accelerations to which each of N charges is subjected, using the $1/r$ law of force

appropriate to rods. If there are electrodes present in which image charges are induced, a force law taking into account the entire array of images created by a single rod can be used instead of the $1/r$ law. For instance, when charges move in a rectangular box with conducting walls, the law of force involves elliptic functions.

This method was used with success in an electron tube problem by Kooyers and Hull at Litton Industries. Of the order of 100 rods were traced in their mutual (plus applied) fields, and their motion was eventually displayed in a movie. This gave considerable insight into the operation of the tube.

However, when it comes to processing the 10^4 rods that a large computer can store, this method is too slow. It would, in the example mentioned, require the computation of some 10^8 elliptic function values at each step!

The method is, in fact, also too accurate. It treats each rod as a (singular) line charge and one often sees close collisions between two such line charges in the model due to the tremendous forces at close range. A rod is better treated as an extended distribution of charge in which the $1/r$ law is cut off at small radii and the "blows" are softened.

One therefore turns to the alternative of preparing a smoothed record of potentials or field components before using them for the acceleration of individual rods. The potential is related to the charges by Poisson's equation and a good deal of research has gone into the numerical solution of this equation.

The methods for solving Poisson's equation have, in general, been iterative methods. They are based on the "relaxation procedure" in which errors are gradually corrected after an inspired initial guess is made of the field. Historically, this method was first applied to the equation for the stress system in an elastic solid. This equation is similar to Poisson's and the errors which one corrects can be physically interpreted as unbalanced local stresses. Hence, the name "relaxation method". A complete calculation of the charge-flow in an ion or electron gun involves, in this procedure, two kinds of iteration--outer loops in which the orbits are retraced through better and better fields and a set of inner loops in

which, from a given charge distribution, the field is calculated to better and better accuracy using guesses, or previous data, as a start.

For the purpose of resolving the more interesting two-dimensional plasma problems, in which fluctuations seem to play a dominant part, this method had to be discarded in favor of a faster method of direct calculation. This could be done at the cost of restricting severely the shapes of the boundaries that can be taken into account. At first, the direct method was in fact restricted to a rectangular box geometry for the boundary surfaces, later it was possible to introduce electrodes into the model which do not form part of the basic rectangular box. The key to solving Poisson's equation by a direct method was Fourier analysis. The rectangular box was subdivided (as in the above mentioned relaxation schemes) into a fine mesh and records were compiled of how many rods, i.e., how many units of charge, fall into each small square. Typically, we have used 48 by 48 mesh points in a large square box of plasma, i.e., 2304 mesh points.

The record of charges is now Fourier analyzed across one dimension. Since in Poisson's equation all Fourier components remain uncoupled from each other, one has effectively reduced the step of calculating potentials from charges to a set of one-dimensional problems--as many as there are Fourier components, i.e., as many as there are mesh points across one side of the large rectangle. Eventually, the total potential has to be Fourier-synthesized from the components obtained from the charge distribution. An exact, noniterative potential distribution is thus obtained. Fourier analysis of the charge records and Fourier synthesis of the potential are the most time-consuming operations in this procedure. By very careful planning, and by utilizing all the symmetries of sines and cosines, R. Hockney at Stanford University was able to develop an overall program for solving Poisson's equation about 10 times faster than the best known iterative method. His program will convert the 2304 charge data to the corresponding 2304 potential data for the mesh points within 0.9 sec on a standard IBM 7090.

This really opens the opportunities for step-by-step transient two-dimensional calculations. Orbits of several thousand particles are

upgraded also within a time of the order of a second, and so a two-dimensional plasma model can be traced with tolerable speed.

Figures 19 and 20 are examples of the results of some two-dimensional calculations made by R. Hockney in which an electron beam is shot into a plasma and at first manages to penetrate it smoothly, there being no significant collisions between beam and plasma particles. However, the beam soon begins to agitate the plasma electrons and to create growing oscillating charge fluctuations and electric fields. Eventually these disordered fields destroy the beam before it can reach the far side of the region investigated in the experiment.

BEAM-GENERATED PLASMAS

If an electron beam is shot into a gas it will ionize the gas atoms, provided the beam energy exceeds the ionization energy which is typically of the order of 20 volts. The gas atoms split into electrons and ions and these created particles flow to the walls of whatever vessel contains the gas. There they recombine at a rate equal to the rate of creation, in a steady-state discharge. A plasma can be created in this manner with a density that is thousands of times the particle density of the beam. This is exactly the process that takes place in an ordinary dc discharge, with the cathode fall region providing the accelerating potential for the beam. It has also been a subject of current interest with kilovolt electron beams being injected from sources external to the plasma.

The same two-stream instability discussed in a previous section can occur in a beam-generated plasma between the electrons of the beam and the electrons of the plasma generated by the beam. It is expected that this instability will be the dominant mechanism for stopping the beam. By the same token, the plasma electrons are heated by the beam, and there is also the possibility of heating the ions by the same mechanism. It therefore appears that this may be a possible method for supplying energy to a thermonuclear plasma, as a starting-torch so to speak. It may also have important applications in the area of microwave power generation.




FIG. 19. Electron trajectories shown in two-dimensional space for an electron beam shot into a plasma. The passage of one unit of time is indicated by the distance between the dots. The beam electrons travel downward and at this point in time reach the wall which is a conducting plane at $Y = 0$ opposite the injection plane at $Y = 48$ which is also a conducting plane. The plasma electrons that are shown move only a small distance near the initial position





FIG. 20. The same problem at a later time. The beam has been stopped by the plasma and the plasma electrons have been violently disturbed.

As in the other problems we have discussed there is a static theory for a plasma formed in this manner, but not including the effect of the beam. The static theory in this case assumes that somehow the plasma electrons have managed to get themselves into a random Maxwellian velocity distribution centered around zero, although they are created with initial velocities in a relatively narrow range centered around a velocity of the order of the ionization potential.

However, the steady-state result need not be a static one, nor a true equilibrium condition. The computer experiment technique seems to be well suited to this problem and work in progress at Stanford by one of the authors and A. S. Halsted is aimed at studying the effects of creating charges with non-Maxwellian velocity distributions and also the effects of beam interactions in a plasma created by the beam. A rather idealized one-dimensional model has been studied in which electrons and ions are created uniformly throughout a diode with zero initial ion velocities and with electrons created all at the same initial speed but with a randomly chosen direction either to right or left. Figure 21 shows the potential in the diode as a function of distance for a number of different times, as obtained from the computer. Also indicated is the static theory prediction. As in the other cases we have discussed, a time-varying potential is obtained, the oscillations occurring in two frequency ranges, one near the electron plasma frequency and one near the ion plasma frequency. Figure 22 shows the total number of charge sheets in the diode as a function of time along with the potential at particular positions in the diode as a function of time. It will be seen that the total number of plasma particles rapidly reaches a nearly constant value but within the diode there are continuing fluctuations.

This example indicates one of a number of new directions being taken in this field; here the new element is that of continuous charge creation within the region of interest.



POTENTIAL PROFILE ACROSS DIODE

FIG. 21. Potential vs. distance in a plasma diode in which electron and ion pairs are created uniformly throughout the diode, all of the electrons being created with the same speed but with a randomly directed velocity to left or right. Ions are created with zero initial velocity and the system is allowed to randomize. The static theory for a perfectly random electron velocity distribution predicts the potential distribution shown.




FIG. 22. The potential as a function of time at two fixed positions in the diode of Fig. 21 and the total number of charge sheets in the diode as a function of time. The total number of charges fluctuates hardly at all, but there are small residual oscillations in the potential near the electron and ion plasma frequencies that are not altered by making the model finer grained.

CONCLUDING REMARKS

After surveying only superficially the variety of computer experiments that have been carried out hitherto, one gets the impression that we are at the threshold of a new era of research. Almost every time that a computer experiment has been carried out seriously it has yielded surprising and significant answers. Almost invariably the experiments have provided deeper insight. Often, as in the case of the thermionic diode, a qualitative or even an analytical quantitative theory has been deduced from one or several computer experiments which allowed one to guess what are the significant effects and what is the correct way of looking at a problem. At times, such enlightened deductions from the first few runs in a series rendered unnecessary all further runs which could instead be predicted by inspired extrapolation.

It is clear that computer experiments will grow in number and importance, not only because of their intrinsic value and success, but also because computers will get bigger and faster, allowing more complex simulations to be made.

One encounters, at times, a prejudice against computer experiments. Partly, such prejudice is based on mathematical snobbery (the formal description of the skin effect in Bessel-functions of complex argument enjoys higher prestige than a few graphs showing how it actually goes!). But often one hears the complaint that a computer can at best say "this is how it happens" and never "this is why it happens". The examples produced here should suffice to answer this complaint. The mere fact that the computer was able to produce the "how" has, many times, told us the "why". Moreover, some users of devices--such as ion engines--do not want to know why a device works, they just want to be sure that it does work. Finally, if a computer merely reproduces and executes the laws of Newton and Maxwell, then it has done its job of proving that these laws are the reason why and that there is no further mystery to worry about.

So far, computer experiments have successfully confirmed actual physical experiments and observations in terms of known basic laws.

Surprises have only come in the form of events that our inability to cope with over-complex situations had prevented us from foreseeing.

A real discovery would be an unsuccessful computer experiment. When all the known interactions and initial data fed into computers fail to reproduce the physical observations, then we shall have to ask the computer to tamper with the laws of nature!

REFERENCES

1. D. A. Dunn and R. E. Holaday, "Electron Trajectories in a Nonuniform Axially Symmetric Magnetic Field," J. Appl. Phys., 32, 1961, pp. 1612-1620.
2. D. R. Hartree, "Some Calculations of Transients in an Electronic Valve," Appl. Sci. Res., B1, 1950, pp. 379-390.
3. R. J. Lomax, "Transient Space Charge Flow," J. Elec. and Control, 9, 1960, pp. 127-140.
4. C. K. Birdsall and W. B. Bridges, "Space Charge Instabilities in Electron Diodes and Plasma Converters," J. Appl. Phys., 32, 1961, pp. 2611-2618.
5. J. Dawson, "One-Dimensional Plasma Model," Physics of Fluids, 5, 1962, pp. 445-458.
6. P. Burger, "Nonexistence of dc States in Low-Pressure Thermionic Converters," J. Appl. Phys., 35, 1964, p. 3048.
7. O. Buneman, "Dissipation of Currents in Ionized Media," Phys. Rev., 115, 1959, pp. 503-517.
8. D. A. Dunn and I. T. Ho, "Computer Experiments on Ion Beam Neutralization with Initially Cold Electrons," Amer. Inst. of Aero. and Astro. J., 1, 1963, pp. 2770-2777.
9. O. Buneman and G. Kooyers, "Computer Simulation of the Electron Mixing Mechanism in Ion Propulsion," Amer. Inst. of Aero. and Astro. J., 1, 1963, pp. 2525-2528.